

# Asymptotically Optimal Gathering on a Grid

Andreas Cord-Landwehr      Matthias Fischer      Daniel Jung  
Friedhelm Meyer auf der Heide

Heinz Nixdorf Institute & Computer Science Department  
University of Paderborn (Germany)  
Fürstenallee 11, 33102 Paderborn

`{cola, mafi, daniel.jung, fmadh}@uni-paderborn.de`

## Abstract

In this paper, we solve the local gathering problem of a swarm of  $n$  indistinguishable, point-shaped robots on a two dimensional grid in asymptotically optimal time  $\mathcal{O}(n)$  in the fully synchronous  $\mathcal{FSYNC}$  time model. Given an arbitrarily distributed (yet connected) swarm of robots, the gathering problem on the grid is to locate all robots within a  $2 \times 2$ -sized area that is not known beforehand. Two robots are connected if they are vertical or horizontal neighbors on the grid. The locality constraint means that no global control, no compass, no global communication and only local vision is available; hence, a robot can only see its grid neighbors up to a constant  $L_1$ -distance, which also limits its movements. A robot can move to one of its eight neighboring grid cells and if two or more robots move to the same location they are *merged* to be only one robot. The locality constraint is the significant challenging issue here, since robot movements must not harm the (only globally checkable) swarm connectivity. For solving the gathering problem, we provide a synchronous algorithm – executed by every robot – which ensures that robots merge without breaking the swarm connectivity. In our model, robots can obtain a special state, which marks such a robot to be performing specific connectivity preserving movements in order to allow later merge operations of the swarm. Compared to the grid, for gathering in the Euclidean plane for the same robot and time model the best known upper bound is  $\mathcal{O}(n^2)$  [DKL<sup>+</sup>11].

**Keywords:** gathering problem, autonomous robots, distributed algorithms, local algorithms, mobile agents, runtime bound, swarm formation problems

## 1 Introduction

In the field of robot formation problems, there is a strong interest in what is feasible when robots only have very limited capabilities. Runtime bounds are intrinsically hard for such problems and tight results are rarely provided. One common formation problem is the gathering, where robots have to gather in one, beforehand unspecified point.

There are essentially three significant approaches that try to tackle this problem: The first approach [DKL<sup>+</sup>11] is a gathering algorithm for  $n$  autonomous, indistinguishable and point-shaped connected robots in the Euclidean plane. A local robot model is used, where the robots especially only have a local vision and no compass. The gathering algorithm works in the  $\mathcal{FSYNC}$  time model and runs in time  $\mathcal{O}(n^2)$ . It is still unknown whether the model-specific time bound is tight

for the general case, too. The second approach [SN14] deals with gathering on a grid in the  $\mathcal{ASYNC}$  time model. The robots have global vision and can move unbounded (finite) steps. The gathering strategy is optimal concerning the total number of movements. The third approach [ACLF<sup>+</sup>16] (accepted for IPDPS 2016, preliminary version [ACLF<sup>+</sup>15]) is an asymptotically optimal gathering strategy for robots, connected as a closed chain on a grid, which is linear in the number of robots in the  $\mathcal{FSYNC}$  time model. Like the first approach [DKL<sup>+</sup>11], the robot model is local, especially only local vision and no compass.

In this work, we use an idea from our gathering algorithm for a closed chain [ACLF<sup>+</sup>16], yet drop the chain connectivity for sake of solving the general gathering on a grid under the same robot and time model. Our solution for the general gathering problem on the grid is a distributed synchronous algorithm that solves the problem asymptotically optimal in the total number of rounds. Note that the  $\mathcal{FSYNC}$  time model is (like in [ACLF<sup>+</sup>16]) one of the main parameters that makes the construction and the analysis of an algorithm very challenging. The combination of preventing symmetry issues due the distributed robot behavior and constant length rounds require the presented complex solution. Contrary, if one would assume a fair scheduler in the  $\mathcal{ASYNC}$  time model, which allows only one robot to be active at a time and finishes a round after every robot has been active at least once, a simple strategy could achieve the same  $\mathcal{O}(n)$  rounds. Specifically when compared to gathering on the Euclidean plane in the  $\mathcal{FSYNC}$  time model, our runtime of  $\mathcal{O}(n)$  is a noteworthy result and beats the best known algorithm, which requires time  $\mathcal{O}(n^2)$  for gathering a swarm of  $n$  robots.

**Our Local Grid Model** Our mobile robots need only simple capabilities: A robot moves on a two-dimensional grid and can change its position to one of its eight horizontal, vertical or diagonal neighboring grid points. A robot can see other robots only within a constant *viewing radius* (measured in  $L_1$ -distance). We call the range of visible robots *viewing range*. The robots controlled by our algorithm need  $L_1$ -distance 20, which can still be optimized. The robots have no compass, no global control, no IDs and no global communication. A robot has a fixed small amount of memory to store a constant number of states. A robot can see the states of all robots inside the viewing range.

Our algorithm uses the fully synchronous time model  $\mathcal{FSYNC}$ . Time is subdivided into equally sized rounds of constant lengths. In every round all robots simultaneously execute their operations in the common *look-compute-move* model [CP04] which divides one operation into three steps. In the *look step* the robot gets a snapshot of the current scenario from its own perspective, restricted to its constant sized viewing range. During the *compute step*, the robot computes its action, and eventually performs it in the *move step*.

**Outline of the Algorithm** The starting point for our problem is our strategy for the gathering of an arbitrary closed chain of robots on a grid [ACLF<sup>+</sup>16]. All operations to shorten the chain based on the strong definition of the robot’s connectivity. Each robot of a chain has exactly two well defined neighbors given by the input of the problem.

The main difference to the general gathering problem is that the connectivity of the chain is missing and so the algorithm could not be applied. This makes general gathering more challenging than in the closed chain case. The idea of this paper is to find a substitute for the chain and the chain’s connectivity in order to transmit the ideas of the algorithm.

The substitute of the chain’s connectivity is the following: We call two robots *connected* if they are horizontal or vertical neighbors on the grid. This means that a robot is connected to at least one and at most four robots. Initially, under the restriction that the swarm is connected, the robots

are arbitrarily distributed on the grid. A swarm is connected, if the robots cannot be separated into two subsets such that no robot of the one subset is connected to any robot of the other subset and vice versa.

In order to find a substitute for the chain we define the so called boundaries: The boundaries consist of all robots who have at least one unconnected side. The black robots of Figure 1 are part of the *outer boundary* and the hatched robots are part of *inner boundaries*. By that definition the swarm consists of one single outer boundary and typically multiple inner boundaries. A robot can detect if it is located on some boundary of the swarm, but because of its limited viewing range it does not know if it is located on the outer boundary or on some inner boundary.




Figure 1: The hatched robots are part of *inner boundaries* and the black robots are part of the *outer boundary*.

The outer boundary has the role as a closed chain and the goal of our algorithm is to shorten the outer boundary. However, the outer boundary does not have the same structure as the chain because the outer boundary can have multiple fringes. Thus our algorithm has to tackle two problems: First, we have to ensure that the algorithm will shorten the outer boundary. Second, we have to ensure that the shortening of the inner boundaries does not disturb the shortening of the outer boundary. Because the robots can not distinguish between inner and outer boundary our algorithm will also shorten the inner boundary.

The main idea for our algorithm is the following: Our algorithm achieves progress in the gathering by performing *merge operations*. A merge operation is a *hop* of a robot located on the boundary onto the same grid cell as one of its neighbors and we remove one of them. A merge operation is only allowed if it does not disconnect the swarm. Sometimes no merge is possible without disconnecting the swarm.

In that case we want the robots of the boundary to reshape the swarm. Reshapement means that the robots reach positions such that new merges are possible, which does not break the connectivity. Reshapement operations consist of two types: Starting a reshaping and continuing a reshaping.

Starting a reshaping is allowed only to a robot that is located on a boundary and can see a certain configuration of robot positions in its viewing range. The robots that are allowed to start a reshaping generate a so called *run state*, hop to another position, and give the run state

to neighbors on the boundary. Robots receiving the run state are called *runners* and are allowed to continue the reshaping by hopping in following rounds and giving the state to neighbors of the boundary. In that way the run state will be moved along the boundary by the runners. If two runners converge and see each other inside their viewing range a merge can be performed. Gathering is finished when all robots are located within a  $2 \times 2$  square, since in our model that situation cannot be simplified anymore. In order to simplify the description of our algorithm, we sometimes allow a robot to reconfigure the positions of the robots in its viewing range.

For counting the progress of our algorithm, it is sufficient to consider the merges of the outer boundary only, i.e. we look at the shortening of the outer boundary. After a constant number of rounds either a merge is possible or a new run state will be generated. A new generated run state will provide at least two runners that converge and enable a merge after  $\mathcal{O}(n)$  rounds. The runners move in parallel and our strategy ensures that the runners do not disturb each other. Different started run states result in different merges. Therefore the algorithm gathers a swarm of  $n$  robots in time  $\mathcal{O}(n)$ . Our result is asymptotically optimal for worst-case swarms.

## 2 Related work

There is a vast literature on robot problems, researching how specific coordination problems can be solved by a swarm of robots given a certain limited set of abilities. The robots are usually point-shaped (hence collisions are neglected) and positioned in the Euclidean plane. They can be equipped with a memory or are *oblivious*, i.e., the robots do not remember anything from the past and perform their actions only on their current views. If robots are anonymous, they do not carry any IDs and cannot be distinguished by their neighbors. Another type of constraint is the compass model: If all robots have the same coordinate system, some tasks are easier to solve than if all robots' coordinate systems are distorted. In [KTI<sup>+</sup>07, ISK<sup>+</sup>12] a classification of these two and also of dynamically changing compass models, as well as their effects regarding the gathering problem in the Euclidean plane, is considered. The operation of a robot is considered in the *look-compute-move model* [CP04]. How the steps of several robots are aligned is given by the *time model*, which can range from an asynchronous  $\mathcal{ASYNC}$  model (for example, see [CP04]), where even the single steps of the robots' steps may be interleaved, to a fully synchronous  $\mathcal{FSYNC}$  model (for example, see [ASY95]), where all steps are performed simultaneously. A collection of recent algorithmic results concerning distributed solving of basic problems like gathering and pattern formation, using robots with very limited capabilities, can be found in [FPS12].

One basic robot formation problem are the shortening and maintenance of a communication chain between two fixed endpoints. The robots then only have a local vision, no compass and a static connectivity is defined by neighborhoods along the chain. For this several results have been published in the  $\mathcal{FSYNC}$  time model. The first shown runtime bound was  $\mathcal{O}(n^2 \log(n))$  [DKLH06]. Later, this has been improved [KM09]: In the Euclidean plane, the *Hopper* strategy delivers a  $\sqrt{2}$ -approximation of the shortest communication chain in time  $\mathcal{O}(n)$ . Restricted to a grid, the *Manhattan Hopper* strategy delivers an optimal solution in time  $\mathcal{O}(n)$ .

One of the most natural problems is to gather a swarm of robots in a single point. Usually, the swarm consists of point-shaped, oblivious, and anonymous robots. The problem is widely studied in the Euclidean plane. Having point-shaped robots, collisions are understood as merges/fusions of robots and interpreted as gathering progress [DKM10]. In [CFPS03] the first gathering algorithm for the  $\mathcal{ASYNC}$  time model with multiplicity detection (i.e., a robot can detect if other robots are also located at its own position) and global views is provided. Gathering in the local setting was studied in [ASY95]. In [Pre07] situations when no gathering is possible are studied. The

question of gathering on graphs instead of gathering in the plane was considered in [Mar09, DFKP06, KMP08]. In [SN13] the authors assume global vision, the  $\mathcal{ASYNC}$  time model and furthermore allow unbounded (finite) movements. They show optimal bounds concerning the number of robot movements for special graph topologies like trees and rings.

Concerning the gathering on grids, in [DDSKN12] it is shown that multiplicity detection is not needed and the authors further provide a characterization of solvable gathering configurations on finite grids. In [SN14], these results are extended to infinite grids, assuming global vision. The authors characterize *gatherable* grid configurations concerning exact gathering in a single point. Under their robot model and the  $\mathcal{ASYNC}$  time model, the authors present an algorithm which gathers *gatherable* configurations optimally concerning the total number of movements.

Assuming only local capabilities of the robots, esp. only local vision and no compass, makes gathering challenging. For example, a given global vision or alternatively just the knowledge of a global compass, the robots could compute the center of the globally smallest enclosing square or circle and just move to this point (global vision) or all robots without any local neighbors in front of them could simply move for example to the south-eastern direction and would finally meet (global compass).

In the  $\mathcal{FSYNC}$  time model, the total running time is a quality measurement of an algorithm. In this time model and additionally under the restriction, that the robots do not have a compass and have only local vision, i.e., they can only see other robots up to a constant distance instead of a global vision of the whole scenery, exist several results that prove runtime bounds. One of these are the strategies for shortening communication chains, we introduced above. I.e., [DKLH06] that needs time  $\mathcal{O}(n^2 \log(n))$  and [KM09] that needs time  $\mathcal{O}(n)$  in the Euclidean plane as well as on a grid.

Our recent result gathers a closed chain on a grid asymptotically optimal in linear time  $\mathcal{O}(n)$  [ACLF<sup>+</sup>16], [ACLF<sup>+</sup>15]. There, connectivity is statically given by a chain structure. For the more general gathering, i.e., using connectivity dynamically given by local vision in the Euclidean plane, an  $\mathcal{O}(n^2)$  runtime bound has been shown [DKL<sup>+</sup>11]. There, every robot synchronously computes the smallest enclosing circle only of the robots within its restricted viewing range and then moves towards its center. Repeatingly executing this synchronous behaviour finally solves the gathering. The authors also prove that for their algorithm the  $\mathcal{O}(n^2)$  bound is tight. For the problem itself, under this local model, a tight bound for the running time is still unknown.

In the present paper, we tighten this bound to the asymptotically optimal value of  $\mathcal{O}(n)$  using the same time- and robot model, but for robots on a grid.

### 3 The algorithm

In our model, merges/removal of robots mean progress of the gathering. Then, at the latest after  $n - 1$  removals the gathering is done. In our model, a robot can be removed, if two robots are located at the same grid cell. Then, one of them is removed.

Because a robot's viewing range is restricted to just a constant size, these removals cannot be performed easily in general. Our algorithm performs two basic operations. More precisely, we have to deal with two cases:

1. Some neighboring robots perform a single hop such that afterwards at least two robots are located at the same position, while preserving the swarm connectivity. Then, we remove one of them. This is a so called *merge* and further discussed in Subsection 3.1.
2. If on some parts of the swarm merges are impossible, we perform so called reshapements to

reshape the swarm in order to prepare merges in succeeding steps. They are explained in Subsection 3.2.

In Subsection 3.1, we now start with case 1. For the sake of simpler descriptions, we use terms like *horizontal*, *vertical*, *downwards*, *left*, *...*. Since our robots do not have a common sense for this, the descriptions/figures are also to be understood in a mirrored or rotated manner.

### 3.1 Merges

A *merge-operation*, or simpler a *merge*, is the simultaneous local operation of a sequence of neighboring robots that merges at least one robot, while not harming the overall connectivity of the swarm. There are different merge operations, all parametrized by a parameter  $k$ , as depicted in Figure 2. In the figure, the case  $k = 1$  denotes the simplest variant, where only a single robot hops onto a grid cell occupied by another robot. Because the robot's viewing radius is  $> 1$ , namely 20, also bigger merges are possible within the local vision. Then, for not to breaking the swarm's connectivity, multiple robots (the black ones in the figure) have to hop simultaneously: Using the local information about the occupied and empty cells in a viewing range, every robot can decide individually, based on its local knowledge, if it participates in the simultaneous merge operation or not. Specifically, the marked white grid cells must not contain any robot, any further marked cell must contain a robot, and the maximal size  $k$  of a merge configuration is limited by the viewing radius. Moreover, all not explicitly depicted cells are ignored for the decision.

The  $k$  robots that constitute a merge form a subboundary of the swarm. When operating, these subboundary robots simultaneously hop one grid cell in the same direction (in the illustration, this means downwards). The swarm's connectivity is ensured, since a merge is only performed when the depicted white cells are empty. By requiring at least one grey cell, i.e., a robot that does not move, after the hops of the subboundary robots at least one robot from a grey cell will be located at the same cell as a robot from a formerly black cell and hence one robot is merged.




Figure 2: Robot subboundaries that allow progress hops (*merge operations*). The value of the *length*  $k$  indicates the number of robots (the black ones) of the subboundary.  $k$  is upper bounded by a robot's constant viewing radius. The operations are only performed, if the cells, marked by white squares, are empty. Else, the swarm's connectivity might break.

Possibly, several simultaneous merges can occur at different parts of the swarm. The problematic cases we have to discuss are those, where two of the robot subsets participating in the merge operations, i.e., the subboundaries, consisting of the black and grey robots of Figure 2, overlap (cf. Figure 3). Precisely, two cases need a closer look (We refer to the black and grey robots of Figure 2.):

1. The beginnings and endings of the subboundaries overlap by two robots.
2. The beginnings and endings of the subboundaries overlap by three robots.

An example for 1) can be seen in Figure 3.a): The robots of the subboundaries 1,2,3 all perform the hop as the black ones during the merge operation (Figure 2). The difference is that afterwards for example the robot *a* is not located at the same position as *b* and vice versa. So, no robot can be removed. But the outermost grey robots of Figure 3.a) do not move and there a merge occurs.

An example for 2) is shown in Figure 3.b). Here, the robot *r* belongs to the both subboundaries 1 and 2. I.e., concerning subboundary 1 *r* would hop downwards, but concerning subboundary 2 it would hop to the left. In this case, *r* performs a diagonal hop to the lower left, while the other robots perform their usual hops. Afterwards, *r*, *a*, *b* occupy the same grid cell and *a*, *b* are removed without breaking the connectivity.




Figure 3: Overlapping merges. (significant examples)

### 3.2 Reshapement of the Swarm by Runners

From the global perspective, if globally nowhere in the swarm the algorithm's local merge operation is possible, then we call it a *Mergeless Swarm*. In this case, our goal is to perform certain *reshapements* of the swarm's outer boundary in order to make merge operations possible. Still considering the swarm from a global perspective, if there is a subboundary as depicted in Figure 4 by the black robots, then we continuously let an outermost robot of this subboundary perform the depicted diagonal hops. Eventually, this will shorten the subboundary enough to allow a (local) merge operation (cf. Figure 2). We say, these hops *reshape* the subboundary.

For an individual robot with its limited viewing range this raises some challenges.

1. When a robot decides to start the reshapement, within its restricted viewing range it does not know if its task of performing diagonal hops will lead to merges (In Figure 4, this is the left outermost black one in round *i*.).
2. In the following rounds, due to the local viewing ranges and skewed coordinate systems, the local view of an outermost black robot may be the same as that of its grey neighbor. So, we



have to ensure that the reshaping is continued by the outermost black robots instead of by their grey neighbors.

3. When robots decide to start the reshaping, symmetries of the subboundary's local shape may lead to breaking of the swarm's connectivity. In Figure 5 this is the case if both  $r$  and  $r'$  start the reshaping.




Figure 4: If the length of the black subboundary (i.e., the value of  $k$  in Figure 2) is larger than the robots' viewing radius, we shrink it by letting one or both of the outermost black robots perform diagonal hops. We indicate such hops by diagonal arrows.




Figure 5: If  $r$  and  $r'$  both start reshaping the subboundary, the connectivity might break.

In the following, we tackle these challenges by introducing certain (local) states of a robot, which we call the *run* state. We call a robot with an active run state a *runner*. Robots can achieve this state in two different ways:

**start runstate:** If the local subboundary within a robot's viewing range has a certain configuration, i.e., the relative positions of other robots and empty cells, then the robot decides on its own to generate the run state. We say, such a robot starts a run. Based on the configuration of the local subboundary, the run state gets a fixed moving direction along the boundary. A robot can start and store up to two run states at the same time. Figure 7 shows how the local configurations must look like.

**move runstate:** A runner  $R(S)$  can move the run state  $S$  to its boundary neighbor  $r'$  in moving direction of  $S$ . We say, the run state has moved from  $R(S)$  to  $r'$ , while its in "start runstate" initially set moving direction always remains unchanged. Afterwards,  $r'$  is identified by  $R(S)$ .



Once a run state has been started in “start runstate”, “move runstate” is executed in every of the following rounds. This means that the run moves along the boundary at constant speed and in the initially settled moving direction.

When starting runs, the shapes of Start-A and Start-B in Figure 7 ensure that the run starts cannot break the swarm’s connectivity. For this, in a situation like in Figure 5 we do not start any runs. We name subboundaries, consisting of shapes like in the figure, *quasi lines*. Figure 6 gives an example of a quasi line.




Figure 6: Example of a quasi line. The fat robots are its endpoints.

**Definition 1** (quasi line). *We call a subboundary a horizontal quasi line, if the following points hold:*

1. *At least its first and last three robots are horizontally aligned.*
2. *All its subboundaries of horizontally aligned robots contain at least three robots.*
3. *All its subboundaries of vertically aligned robots contain at most two robots.*

*In a Mergeless Swarm, at both ends of a quasi line a run starting subboundary Start-A or Start-B of Figure 7 in a matching rotation or reflection occurs. (If the swarm is not mergeless, then the subboundaries outside the quasi line’s endpoints may also have other shapes than these.)*

*The definition of a vertical quasi line follows analogously.*

We let runs start at endpoints of quasi lines. We let them move along such quasi lines and, while doing this, perform reshapements of the boundary. For this, depending on the local shape of the subboundary, we require the run operations, shown in Figure 8.

- a) OP-A: The runner and at least the next 3 robots are located on a straight line. Here, the runner first performs a diagonal hop, then moves the run to the next robot.
- b) OP-B: The runner and only the next 2 robots are located on a straight line. Then, for 3 times the runners just move the run to the next robot without any diagonal hops. Afterwards, it is located at the target corner  $c$ .
- c) OP-C: This one is needed at most once for a new run, if started at the run starting subboundary Start-B (Figure 7.(ii)).

These operations shall finally enable merges.

Two properties let multiple runs be active in parallel.

1. Usually, at different positions, the whole swarms contains multiple run starting subboundaries (cf. Figure 7) at the same point in time.

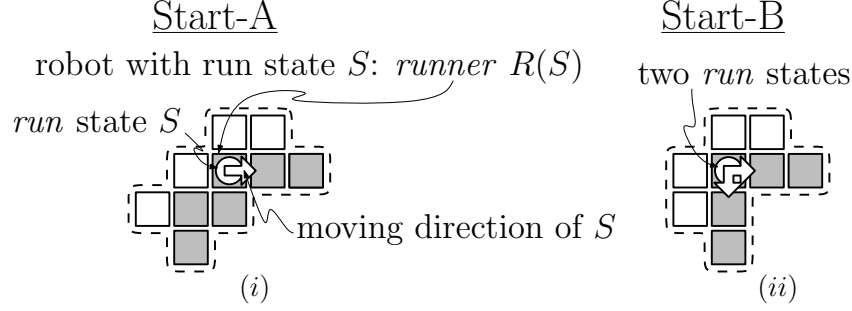


Figure 7: Run starting subboundaries: The robots marked by white circles decide to start the run states, only based on the relative positions of the marked robots and empty cells. Here, grey squares denote robots, while white squares denote cells that must be empty. The white arrows indicate the moving direction of the runs. This is the notation, we will use for marking a runner.  $R(S)$  identifies the runner/robot which currently has the run state  $S$ ; In (ii), Start-B, the robot marked by the circle is the endpoint of a horizontally and a vertically aligned subboundary at the same time. Here, we must start two runs, moving in both directions along the boundary.

2. No matter, if previously started runs are still active, every constant number of  $L = 22$  rounds all robots simultaneously check if they can start new runs (cf. Figure 7) and if so, they do so.

So, we have to deal with two runs, meeting on a quasi line. Then, we distinguish two cases (Cf. Figure 9.). a) This run pair has enabled the desired merge. b) This pair is oriented in a way that does not enable a merge. We let the runs of such pairs pass along each other. Figure 9.b) shows for the case that both runs are located on the same quasi line how this is performed: At the time when their distance along the boundary (i.e., the number of robots on the subboundary connecting both  $+1$ ) is 3 or less, they only keep moving along the boundary, but the runners do not perform reshaping hops. We call 3 the *run passing distance*. This is repeated until  $S_1$  is located at its target robot  $c_2$  (then also  $S_2$  is located at its target  $c_1$ ). We call this the *run passing operation*. Afterwards, the normal reshaping operations are continued. In Section 6, we provide a full explanation of all possible situations of needed run passing operations. While in the easy case a viewing radius of 11 is the minimum value that suffices, we then need the unoptimized value of 20.

### 3.3 Stopping Runs

We use the example of Figure 10 for introducing some terms when speaking about runs: Because  $S_1, S_2$  are moving in the same direction, we call them *sequent* runs, while relative to their moving direction  $S_1$  is located *in front of*  $S_2$ . The *distance* between them is defined as the number of robots on the subboundary connecting both  $+1$ . We say,  $S_1$  is *visible* to  $S_2$ , if the *distance* between them is  $\leq 20$ .

In order to ensure that runs, being active at the same time, can work correctly, some of them have to stop in certain situations. More precisely, we let a runner  $R(S)$  stop/terminate its run  $S$ , if one of the conditions of Table 1 is true.

We give some more detailed explanations concerning some of the conditions of Table 1:

- 1) For example, because of merges two sequent runs may come too close to each other, which then might hinder the pipelining. The affected runners can detect this on their own. The




Figure 8: a) OP-A: Reshapement by a runner. The operation takes only one round. b) OP-B: No diagonal hops are performed until the target corner  $c$  is reached. c) OP-C: Special case when new runs start: First, perform one diagonal hop, then no diagonal hops until the target corner  $c$  is reached.

criterion for this is that the next sequent run in front of them becomes visible. Then, the termination condition 1) matches and the run behind stops.

- 4) We assume, that in Figure 9.b) to the right of robot  $c_2$  another run  $S_3$ , moving in the same direction as  $S_2$ , is located. During the rounds in which  $S_1$  and  $S_2$  are performing their run passing operation,  $S_3$  keeps moving towards  $c_2$ . Now, for example, it may happen that because of the reshapements of  $S_3$ ,  $c_2$  becomes part of a merge operation. Then,  $c_2$  would hop downwards such that the corner shape does not exist anymore. Because this corner has been the target of  $S_1$ ,  $S_1$  could not continue its reshapements after the run passing, so it terminates.

Summarizing, all robots synchronously execute the algorithm, shown in Figure 11.

## 4 Why the strategy produces progress in gathering

For measuring progress, we only consider runs, started pairwise at both ends of certain quasi lines. The kind of runs pairs which in any case lead to a merge, are called *good pairs*.




Figure 9: a) *Merge operation*: The reshapements of the runs  $S_1, S_2$  have enabled a merge. This merge is performed and both are stopped. The operation takes only one round. b) *Run passing operation*: The runs  $S_1, S_2$  cannot enable a merge. If their distance along the boundary is less or equal 3, then they pass each other by only keeping moving but without making the runners perform diagonal hops. Afterwards, i.e., when  $S_1, S_2$  have reached their target robots/corners  $c_2, c_1$ , they return to normal operation.




Figure 10: The *distance* between  $S_1, S_2$  is 8.

#### 4.1 Good pairs

Assume a newly started run pair, connected by a quasi line. We call this pair a *good pair*, if the following is true (cf. Figure 12): The exterior neighbors (the fat bordered robots in the figure) of the run pair are both located downwards and the whole (global) area above the connecting quasi line does not contain any robot (The same definition analogously holds for any rotation.). Looking at our run operations OP-A, OP-B and OP-C as well as the run passing operation and merges, we notice that none of them could let a robot which is not part of the considered quasi line hop into this empty area.

Such good pairs always enable a merge if they have been moving close enough together and then terminate. Figure 13 shows an example of a good pair where the shape of the quasi line connecting both, is a straight line. The example of the figure shows that using runs, we solve the problem, we have had in Figure 4, where it was impossible to locally decide by the robots, which of them has to perform the reshapement hop. Figure 14 shows an example for a more general case of a quasi line: Again, the black subboundary will be reshaped for performing a merge. But now, several run operations (cf. Figure 8) are needed until a runner arrives at an endpoint of the black subboundary: Until round  $i + 4$ , the runners execute only the operation OP-A (cf. Figure 8). Afterwards, the runner  $R(S_r)$  starts operation OP-B, while  $R(S_\ell)$  still executes OP-A. After operation OP-B has been processed completely,  $S_r$  finally is located at the right end of the

A runner stops/terminates its run, if at least one of the following conditions is true:

1. It can see the next sequent run in front of it (This happens if sequent runs have come too close to each other, e.g., because of merge operations.).
2. It can see the quasi line’s endpoint in front of it.
3. It was part of a merge operation.
4. While it performs the run passing operation (Figure 9.b)), some change of the subboundary’s shape prevents the operation from being completed successfully. (This can happen because of a merge.)
5. While it performs the operation OP-B or OP-C (Figure 8), some change of the subboundary’s shape prevents the operation from being completed successfully. (This can happen because of a merge.)
6. While it performs the operation OP-A or OP-C, it has hopped onto an occupied cell. (Then, one of both robots is removed.)

Table 1: Conditions which let a run terminate.

black subboundary (round  $i + 8$ ). Now, its runners start shortening this subboundary by executing OP-A until the merge can be performed. Afterwards,  $S_\ell$  is still active and keeps moving and will stop at the latest when an endpoint of the quasi line becomes visible.

Many of the started runs do not belong to good pairs. In Section 5, we will prove the actual existence of good pairs.

## 4.2 Pipelining

We will show that even if multiple good pairs are nested into each other, different good pairs will enable different merges. This is what we call *pipelining*. Figure 15 shows an example for the case that the quasi line has the shape of a straight line. The more general case works analog to this.

As runs are moving with constant speed, the runs of the inner good pair will meet and as the result enable a merge and stop, first. Then, obviously the outer good pair will also enable a (different) merge, some rounds later.

## 5 Correctness and running time

Our basic argumentation for the total running time is the following: If every round a merge can be performed, then the time needed for the gathering is obviously upper bounded by  $n$ , with  $n$  being the number of robots. If no merge can be performed, then every  $L = 22$  rounds a *good pair* can be started which reshapes the swarm such that a merge can be performed after at most  $n$  further rounds. As  $L$  is a constant, this then leads to a linear total running time  $\mathcal{O}(n)$ .

In this section, we give the formal proofs. Our goal is the proof of the correctness and the linear running time (Theorem 1). In our main approach, we want to enable merges by good pairs if else no merge could be performed. The following two lemmas prove that this actually works. They are the base for the proof of the theorem. The proofs of the lemmas can be found in Subsection 5.1 and 5.2.

Every robot  $r$  every round checks the following three steps:

1. Merge: If  $r$  detects a possible merge within its viewing range then
  - if  $r$  is one of the black robots in Figure 2, it hops downwards.
  - if afterwards  $r$  is located at the same position as one of the grey robots, the grey one is removed without breaking the connectivity.
2. Run Operations: If  $r$  is a *runner*, then
  - (a) Its run terminates/stops if any of the conditions of Table 1 is true.
  - (b) Runner's Movement and Reshapement
    - Run passing:
      - If  $r$  is currently in progress of executing the run passing operation (Figure 9.b)), then this operation is continued.
      - Else, if  $r$  can see a run in front of it, such that both are moving towards each other and the distance between them is less or equal than the *run passing distance*, then  $r$  starts the run passing op.
    - If  $r$  is not in progress of passing, then
      - If  $r$  is in progress of executing a run operation OP-B or OP-C (cf. Figure 8), which take more than one round, then this one is continued.
      - Else:  $r$  executes the matching new run operation OP-A, OP-B or OP-C.
3. Start new runs: Every  $(L = 22)$ th round,  $r$  checks if it can start a new run:  
 If  $r$  is one of robots marked by circles, of the run starting subboundaries Start-A or Start-B (cf. Figure 7), then it starts one resp. two runs.

Figure 11: The algorithm.

**Progress pairs** Because of the local vision, new runs and maybe also good pairs are started every  $L = 22$  rounds, regardless of whether or not at somewhere in the swarm merges can be performed. In the following analysis, we will only argue with new good pairs which are started if during the last  $L - 1$  and the current round in the whole swarm no merge has been performed. We distinguish such good pairs from others by calling them *progress pairs*.

**Lemma 1.** *Every  $L = 22$  rounds either a merge has been performed or else a new progress pair is started.*

**Lemma 2.** *For all progress pairs the following properties hold.*

- a) *Every progress pair enables a merge (after at most  $n$  rounds).*
- b) *Different progress pairs enable different merges.*

Using these two lemmas, we can now prove the total linear running time.

**Theorem 1.** *Given a swarm of  $n$  robots. Then, after  $\mathcal{O}(n)$  rounds gathering is done. This is asymptotically optimal.*




Figure 12: Pair of runs, connected by a quasi line. The runs are a *good pair* if the outer subboundary neighbors (the fat robots) are both located downwards and the globally the whole area above does not contain any robots.




Figure 13: A *good pair* of runs. In round  $i$ , the new runs start (cf. Figure 7). Afterwards, the action OP-A of Figure 8 are repeatedly executed. The runs are moving closer and closer together. If the runs have moved close enough, a merge can be performed.

*Proof.* We subdivide time into intervals of lengths  $L$ , where  $L$  denotes a number of rounds. Merges can be performed during at most  $n$  such intervals, because every merge removes at least one robot. In all other intervals a new progress pair starts (Lemma 1). Each of these progress pairs leads to a merge (Lemma 2.a)). Because no two of them lead to the same merge (Lemma 2.b)), the number of intervals without merges is also upper bounded by  $n$ .

By Lemma 2, a progress pair needs at most  $n$  rounds until it has led to a merge. We assume the worst-case, which is that in the last of the  $2 \cdot n$  intervals the last progress pair was started. Then the total running time is upper bounded by  $2n \cdot L + n$ , which proves the upper bound of the theorem, because  $L$  is a constant.

In our model, the diameter of the initial configuration provides the worst-case lower bound  $\Omega(n)$  for any gathering strategy.  $\square$






Figure 14: A good pair, started on a quasi line in order to shorten the black subboundary. Several and different run operations are needed.

### 5.1 Proof of Lemma 1

*Proof.* The lemma requires that the swarm is a *Mergeless Swarm*. As the basis of our proof, we start by showing that then the outer boundary of the swarm only consists of quasi lines and stairways (cf. Figure 16). First, we assume that every robot is connected to only two other robots. Then, especially no robots are located at the inside of the swarm, and the swarm consists only of its *outer boundary*. We start with the assumption that merges are only possible up to the length 2 (cf. Figure 2). Then, as by definition all horizontal subchains of a horizontal quasi line consist of at least 3 robots, no merge can be performed on them. In order to connect two horizontal or two vertical quasi lines without enabling a merge, they must be connected by so called *stairways*. Stairways can have arbitrary length and are subchains of alternating left and right turns. Figure 16 shows an example. The bicolored robots are the connection points between the quasi lines and the stairway. All differently shaped connecting subchains would allow merges. So, assuming that every robot is connected to only to other robots, the whole outer boundary consists only of quasi lines and stairways. This clearly is also the case, if longer merges are allowed as, for horizontal quasi lines, increase the number of horizontally aligned robots but decrease the number of vertically aligned ones (cf. Definition 1 and Figure 2).

Now, we also allow connectivities bigger than 2. Then robots may also be located at the swarm's inside or some parts of the outer boundary may be neighboring (i.e., if at that position the diameter = 2). Such robots can hinder merges. Figure 17 shows examples for configurations with existing inside robots. In the figure, the dashed bordered robots are part of the outer boundary. The hatched robots prevent the black ones from performing a merge for which the black ones needed to hop upwards. So, although the swarm is mergeless, the outer boundary may still contain mergable shapes. For merges of length  $> 2$  ((i) in the Figure), the outer boundary then still is a quasi line. For merges of length 1 and 2 ((i, ii) in the Figure), the black robots are not part of the outer boundary. So, in every *Mergeless Swarm*, the outer boundary only consists of quasi lines and stairways.

Now, we can show that in a *Mergeless Swarm* always a progress pair exists. For this, we assume a global north and west and construct a vector chain along the swarm's outer boundary as follows.




Figure 15: Pipelining of runs. New runs are started every  $L = 22$  rounds.




Figure 16: Two quasi lines, connected by a *stairway*.

We consider the upper envelope of the swarm and take its left- and rightmost robots  $s$  respectively  $t$ . Because  $s, t$  must be located on the boundary of the swarm's smallest enclosing rectangle and the swarm is mergeless, below each of them must exist two additional robots (cf. Figure 18). We start from the left at the vector  $\vec{v}_0$  constructing a vector chain in clockwise orientation along the outer boundary of the swarm, ending at the vector  $\vec{v}_m$ . The figure shows a significant example of this construction. Here, the grey colored robots denote the corresponding part of the swarm's outer boundary. Starting from its left end, we divide this vector chain into longest  $x$ -monotone subchains. The vector, originated at  $s$ , by construction points to the east. The second subchain starts when the first vector points to the west. The third subchain starts when the first vector points to the east again, and so forth. In the figure, these are the vectors  $\vec{v}_{16}$  respectively  $\vec{v}_{26}$ . The constructed vector chain may overlap itself at places where the diameter of the swarm's boundary amounts only 1, but cannot contain any crossings. Because all constructed subchains are  $x$ -monotone, at least one of them is fully part of the upper envelope. Let  $\mathcal{C}$  denote this subchain. If we consider  $\mathcal{C}$ , then, as the swarm is mergeless, both predecessors of its first vector must point to the north, i.e., the corresponding robots are located downwards. By the same arguments, both successors of its last vector must point to the south, i.e., the corresponding robots are also located downwards. This also holds if  $s$  or  $t$  the endpoints of this vector chain. Then, on both sides are three neighboring robots on a vertical straight line. By Definition 1 of a quasi line, these robots must be part of vertical quasi lines.

In contrast, if looking at the subchain  $\mathcal{C}$ , then because the swarm is mergeless,  $\mathcal{C}$  must contain at least two succeeding vectors, both pointing to the east. The three connected robots then must be part of a horizontal quasi line. At the transition between horizontal and vertical quasi lines,




Figure 17: The hatched robot prevents the merge.




Figure 18: Constructing a vector chain along the swarm's outer boundary.

possibly connected by stairways, the shape matches the starting subboundaries Start-A or Start-B of Figure 7. As both vertical quasi lines are located downwards w.r.t. the horizontal quasi line of  $\mathcal{C}$ , a good pair can be started (cf. Subsection 4.1). Because we have assumed that the swarm has been a Mergeless Swarm also during the previous  $L - 1$  rounds, this good pair is a progress pair.  $\square$

## 5.2 Proof of Lemma 2

For the proof, we need the run invariants of Lemma 3.

**Lemma 3.** *The value  $L = 22$  and the value 20 for the viewing radius ensure, that for every run  $S$  until it terminates, the following invariants holds.*

1. *Every round,  $S$  moves one robot further in moving direction.*
2. *After the first three rounds after its start the reshapements of the runner  $R(S)$  do not violate the quasi line Definition 1 of its own quasi line.*
3. *The reshapements of  $R(S)$  do not violate the quasi line shape of other quasi lines if they belong to good pairs.*
4.  *$S$  cannot see other sequent runs in front of it.*

5.  $S$  is either in progress of passing along another run or the runner  $R(S)$  executes one of the operations  $OP-A$ ,  $OP-B$  or  $OP-C$  (Figure 8).
6. Good pairs stay being good pairs.

Now, we can prove Lemma 2:

*Proof. a)* : At the time, a progress pair  $S, S'$  is started, the subboundary connecting both is a quasi line. Because of Lemma 3.2) and 3), this also does not change if other runs are located on this quasi line or next to it. And also merges preserve the quasi line properties. So, if not stopped,  $S$  and  $S'$  keep moving towards each other (Lemma 3.1)). Because of Lemma 3.6) a merge can be performed at the latest when they meet, which takes at most  $n$  rounds.

It remains to show that  $S, S'$  are not stopped by the algorithm's termination conditions (Table 1) before the merge could be performed. In the following, we check all these conditions. In this proof, we denote the quasi line, on which  $S, S'$  are located by the variable  $q$ . Let  $S^*$  be the next sequent run in front of  $S$ . 1): By definition of a progress pair, no merge has been performed since the last time new runs have been started. This means, that if no run from outside the quasi line  $q$  has interfered with  $q$ , the distance between  $S$  and  $S^*$  is at least  $L - 1$  (cf. the proof of Lemma 3.5)) which is bigger than the viewing radius. (The same analogously holds for  $S'$ .) So then, the runs of a progress pair cannot be stopped by this termination condition. So, we assume that other runs have interfered with  $q$ . If  $S^*$  originates from a different quasi line,  $S$  and  $S^*$  can only be oriented in the way shown in Figure 19, because the run operations require the cells "above" be empty. If these quasi lines shall meet, then at least one of them cannot be  $y$ -monotone. But then a merge would have removed at least one of them before both runs have come too close. So  $S^*$  cannot originate from a different quasi line.




Figure 19: Too close sequent runs cannot originate from different quasi lines.

So, a run on a different quasi line (lying alongside  $q$ ) must have shortened the part of  $q$  that connects  $S$  and  $S^*$ . This can only have happened because of a merge operation, but this is not possible, because  $S, S'$  is a progress pair.

2): If a run of a progress pair can see an endpoint of the quasi line in front of it, then the other run of the progress pair must have previously been stopped (Using the properties of Lemma 3.2) and 3).). Then, because a run of a progress pair cannot be stopped by condition 1), it instead must have been stopped because of a merge. So, a termination because of condition 2) is allowed.

4, 5): The change of the quasi lines shape could either have happened because of the reshaping of a sequent run or because of a merge operation. Because we have ensured the minimum distance of two sequent runs to be large enough, the first case cannot happen. So, a merge must have been the reason. Then, this merge must have been enabled by another run  $S^*$ , moving towards  $S$ . If  $S^*$  is the partner run of  $S$ , then stopping is allowed. Else, because progress pairs are nested into each other,  $S^*$  either was not a run of a progress pair or its partner run has previously been stopped. In the latter case, because runs of progress pairs cannot be stopped because of condition 1), the run must have been stopped by an earlier merge. In both cases, the current merge can be credited to the progress pair of  $S$ .

3): By the same arguments, this merge does not need to be also credited to a different progress pair. This then also proves b) of the lemma.  $\square$

The missing proof of Lemma 3 can be found in Subsection 5.3.

### 5.3 Proof of Lemma 3

*Proof.* Let  $q$  be the quasi line  $S$  is located on and we w.l.o.g. assume that  $q$  is a horizontal quasi line. 1): This directly follows by the run definition in Subsection 3.2.

2): Cf. Figure 8. OP-A ensures that on horizontal quasi lines only horizontal subchains of lengths  $> 2$  are shortened and vertical subchains remain unextended. OP-B does not reshape. OP-C can only be executed during the first three rounds after  $S$  has been started. Our design of the run passing operation also ensures that the quasi line properties are maintained.

3): Because of the run stopping condition of Table 1.2),  $R(S)$  can only reshape a quasi line  $q^*$  if  $q^*$  is also a horizontal one. As the run operation OP-C cannot be applied inside  $q^*$ , we only need to consider OP-A. If the cell onto that  $R(S)$  hops, is empty, then OP-A (because of the cells that are required to be empty) can only be executed if  $S$  is also located on  $q^*$ . But then, by the same arguments as in 2), the reshaping of  $R(S)$  cannot violate the quasi line definition. If in contrast the cell onto the  $R(S)$  hops is occupied by some of the robots of  $q^*$ , then this operation does not modify the shape of  $q^*$ .

4): This is ensured by the run termination condition of Table 1.1).

5): Cf. Figure 8. All run operations OP-A, OP-B and OP-C ensure that if  $S$  was located at some corner  $c_1$  when the operation started, it afterwards either is located at some other corner  $c_2$  such that both corners are rotated equally or terminates if the target corner has been removed during the operation (cf. Table 1.5, 4)). Then, because still located on a quasi line (cf. 2)), again a) or b) can be applied or run passing is started.

The run passing needs some closer look, because it interrupts other operations. In order to ensure a regulated behavior, we chose the distance between sequent runs big enough such that a run does not have to execute a new run passing operation before it has finished its previous one. As the value of the constant  $L$  controls the length of the waiting intervals between the start of two sequent runs and our stopping conditions maintain this minimum distance by stopping a run if the next sequent run in front of it becomes visible, we need to settle the values for these constants appropriately. At this place, we only analyze the case that was explained in Section 3.2, i.e., all participating runs are located on the same quasi line. The other cases (Section 6) can be proven similarly. We look at two sequent runs  $S_1$  and  $S_1^{\text{succ}}$  such that  $S_1^{\text{succ}}$  has been started after  $S_1$ . Their distance  $D$  is at least  $L - 1$ . This value is achieved if at the start of  $S_1$  the operation OP-C (Figure 8) was executed.

Now, we chose the value for the constant  $D$  big enough for ensuring that while some run  $S_2$  is passing along  $S_1$ ,  $S_1^{\text{succ}}$  becomes visible to  $S_2$  the earliest when the passing operation with  $S_1$  has

been completed. Figure 20 shows an example for the longest possible duration of a run passing operation.




Figure 20: Runs, passing along each other while the runner  $R(S_1)$  is executing the operation OP-B (cf. Figure 8).

Here, it takes 6 rounds until  $S_2$  has arrived at its target corner. Because the run passing operation starts when the distance between  $S_1$  and  $S_2$  is  $\leq 3$ , after the passing operation, the distance between  $S_2$  and  $S_1^{\text{succ}}$  equals  $D - 9$ . We want this then still be  $\geq 3$ . So we choose  $D \geq 12$  and together with the above argumentation concerning the minimum distance between sequent runs follows  $L \geq 13$ . In order to detect that the distance has become smaller than 12 (and solve this problem), the viewing radius must be 11 (cf. Table 1.1)).

Using the same construction, but including all other possible cases of run passings (Section 6), we can show that (unoptimized) values of  $L = 22$  a viewing radius of 20 suffice.

6): When defining good pairs in Subsection 3.2, good pairs have been characterized by the relative position of the outer direct neighbors of the good pair according to the quasi line. The first part of the proof of 5) finishes the proof.  $\square$

## 6 Run Passing Operation in Detail

If two runs  $S, S'$  that are moving towards each other but do not enable a merge, have come too close to each other, then we let them pass along each other with a constant movement speed but without performing diagonal hops, in order to prevent the connectivity from breaking or destroying the quasi line property. We call this the *run passing* operation. If they have come this close that the operation is needed, then also  $S$  can see  $S'$  withing its constant viewing range and vice versa. For the run passing operation, it is important to assign a target corner to every such run, so that it can continue its reshapements afterwards. Cf. Figure 21. The subfigures show two runs  $S, S'$  that need to pass along each other. The curve  $q$ , originating at  $S$  symbolizes the quasi line that belongs to  $S$ . Analogously, the quasi line  $q'$  belongs to  $S'$ . We can assume that locally no robots are located in the area above  $q$  and below  $q'$ , because else the corresponding run  $S$  respectively  $S'$  could not be part of a good pair so that it is allowed to stop itself. Then, the run passing is not needed anymore. If a run passing operation is necessary, then, between the locations of  $S, S'$ , the quasi lines  $q, q'$  must in some way overlap. We now explain that in all possible cases of overlappings always suitable target corners can be assigned to both of the runs, so that they can continue their reshapements after the passing. We refer to Figure 21.

- a)  $q, q'$  are identical. Then the run passing operation executes in the way explained in Section 3.2.
- b)  $q, q'$  overlap at the location of  $S$  and at the location of  $S'$ . Then by the same arguments as in a), the target corners  $c, c'$  must exist.




Figure 21: Classification of all possible cases that can occur if two runs need to pass along each other.

- c)  $q, q'$  are disjoint. Then the runners calculate whether their reshapements would make  $q$  and  $q'$  overlap. If so, then this overlapping would lead to the removal/merge of robots, so that we can credit this merge to the good pair and both runs are allowed to stop. Else, the runs continue their normal reshapement operations.
- d)  $q, q'$  overlap only at the location of  $S'$ . Then, the vertical distance between  $S$  and  $S'$  must be  $\geq 0$  because else a merge would make a run passing superfluous. As on the right,  $q'$  is located below  $q$  ( $d \geq 1$ ), there must exist some corner  $c^*$  with the same rotation as the corner at that  $S'$  is located. The robots then reconfigure the robots of  $q'$  this way, that  $c^*$  is located at the right endpoint of the sub quasi line. Afterwards,  $S'$  can move to this corner and then continue its reshapements. In contrast,  $S$  can simply move to  $c'$ .
- e)  $q, q'$  overlap but are disjoint at both endpoints. Then, by the same arguments as in d), both target corners exist.

## References

- [ACLF<sup>+</sup>15] ABSHOFF, Sebastian ; CORD-LANDWEHR, Andreas ; FISCHER, Matthias ; JUNG, Daniel ; HEIDE, Friedhelm Meyer auf d.: Gathering a Closed Chain of Robots on a Grid. In: *CoRR* abs/1510.05454 (2015). <http://arxiv.org/abs/1510.05454>
- [ACLF<sup>+</sup>16] ABSHOFF, Sebastian ; CORD-LANDWEHR, Andreas ; FISCHER, Matthias ; JUNG, Daniel ; HEIDE, Friedhelm Meyer auf d.: Gathering a Closed Chain of Robots on a Grid. In: *IPDPS '16*, IEEE, to appear, Mai 2016
- [ASY95] ANDO, Hideki ; SUZUKI, Yoshinobu ; YAMASHITA, Masafumi: Formation and agreement problems for synchronous mobile robots with limited visibility. In: *ISIC '95*, 1995, S. 453–460
- [CFPS03] CIELIEBAK, Mark ; FLOCCHINI, Paola ; PRENCIPE, Giuseppe ; SANTORO, Nicola: Solving the Robots Gathering Problem. In: *ICALP '03*, 2003, S. 1181–1196
- [CP04] COHEN, Reuven ; PELEG, David: Robot Convergence via Center-of-Gravity Algorithms. In: *SIROCCO '04* Bd. 3104, 2004 (LNCS), S. 79–88



- [DDSKN12] D'ANGELO, Gianlorenzo ; DI STEFANO, Gabriele ; KLASING, Ralf ; NAVARRA, Alfredo: Gathering of Robots on Anonymous Grids without Multiplicity Detection. In: *SIROCCO '12* Bd. 7355. 2012. – ISBN 978–3–642–31103–1, S. 327–338
- [DFKP06] DESSMARK, Anders ; FRAIGNIAUD, Pierre ; KOWALSKI, Dariusz R. ; PELC, Andrzej: Deterministic Rendezvous in Graphs. In: *Algorithmica* 46 (2006), Nr. 1, S. 69–96
- [DKL<sup>+</sup>11] DEGENER, Bastian ; KEMPKES, Barbara ; LANGNER, Tobias ; MEYER AUF DER HEIDE, Friedhelm ; PIETRZYK, Peter ; WATTENHOFER, Roger: A tight runtime bound for synchronous gathering of autonomous robots with limited visibility. In: *SPAA '11*, 2011, S. 139–148
- [DKLH06] DYNIA, Mirosław ; KUTYŁOWSKI, Jarosław ; LOREK, Paweł ; HEIDE, Friedhelm Meyer auf d.: Maintaining Communication Between an Explorer and a Base Station. In: *IFIP TC10*, 2006, S. 137–146
- [DKM10] DEGENER, Bastian ; KEMPKES, Barbara ; MEYER AUF DER HEIDE, Friedhelm: A local  $O(n^2)$  gathering algorithm. In: *SPAA '10*, 2010, S. 217–223
- [FPS12] FLOCCHINI, Paola ; PRENCIPE, Giuseppe ; SANTORO, Nicola: *Distributed Computing by Oblivious Mobile Robots*. Morgan & Claypool, 2012 (Synthesis Lectures on Distributed Computing Theory)
- [ISK<sup>+</sup>12] IZUMI, Taisuke ; SOUISSI, Samia ; KATAYAMA, Yoshiaki ; INUZUKA, Nobuhiro ; DÉFAGO, Xavier ; WADA, Koichi ; YAMASHITA, Masafumi: The Gathering Problem for Two Oblivious Robots with Unreliable Compasses. In: *SICOMP* 41 (2012), Nr. 1, S. 26–46
- [KM09] KUTYŁOWSKI, Jarosław ; MEYER AUF DER HEIDE, Friedhelm: Optimal strategies for maintaining a chain of relays between an explorer and a base camp. In: *TCS* 410 (2009), Nr. 36, S. 3391–3405
- [KMP08] KLASING, Ralf ; MARKOU, Euripides ; PELC, Andrzej: Gathering asynchronous oblivious mobile robots in a ring. In: *TCS* 390 (2008), Nr. 1, S. 27–39
- [KTI<sup>+</sup>07] KATAYAMA, Y ; TOMIDA, Y ; IMAZU, H ; INUZUKA, N ; WADA, Koichi: Dynamic Compass Models and Gathering Algorithms for Autonomous Mobile Robots. In: *SIROCCO '07* Bd. 4474, 2007 (LNCS), S. 274–288
- [Mar09] MARTÍNEZ, Sonia: Practical multiagent rendezvous through modified circumcenter algorithms. In: *Automatica* 45 (2009), Nr. 9, S. 2010–2017
- [Pre07] PRENCIPE, Giuseppe: Impossibility of gathering by a set of autonomous mobile robots. In: *TCS* 384 (2007), Nr. 2-3, S. 222–231
- [SN13] STEFANO, Gabriele D. ; NAVARRA, Alfredo: Optimal Gathering of Oblivious Robots in Anonymous Graphs. In: *LNCS* 8179 (2013), S. 213–224
- [SN14] STEFANO, Gabriele D. ; NAVARRA, Alfredo: Optimal Gathering on Infinite Grids. In: *SSS '14*. 2014, S. 211–225